

TOP500 List's Twice-Yearly Snapshots of World's Fastest Supercomputers Develop Into Big Picture of Changing Technology

BY ERICH STROHMAIER

COMPUTER SCIENTIST, FUTURE TECHNOLOGIES GROUP, LAWRENCE BERKELEY NATIONAL LABORATORY, AND FOUNDING CO-EDITOR, TOP500 LIST

estrohmaier@lbl.gov

Abstract

Now in its 10th year, the TOP500 list of supercomputers serves as a “Who’s Who” in the field of high performance computing (HPC). The TOP500 list was started in 1993, compiling and publishing twice a year a list of the most powerful supercomputers in the world. But it is more than just a ranking system and serves as major source of information for analyzing trends in HPC. In this article we analyze some major trends in HPC based on the quantitative data gathered over the years in this TOP500 project (see www.top500.org for a complete access to all data).

The list of manufacturers active in this market segment has changed continuously and quite dramatically during the 10 year history of this project. And while the architectures of the systems in the list have also seen constant change, it turns out that the overall increase in the performance levels recorded is rather smooth and predictable. The most important single factor for this growth is the increase of processor performance described by Moore’s Law. However, the TOP500 list clearly illustrates that HPC performance has actually outpaced Moore's Law, due to the increasing number of processors in HPC systems.

Introduction

During the 1980s at the University of Mannheim, Germany, we started collecting data and publishing statistics about the supercomputer market. At that time it was relatively simple to define what a supercomputer was, as vector systems such as the Cray Y-MP delivered otherwise unmatched computing performance. Thus, a simple count of vector systems provided good statistics of the HPC market. At the beginning of the 1990s, a considerable number of companies competed in the HPC market with a large variety of architectures, such as vector computer, mini vector computer, SIMD (single instruction on multiple data) and MPP (massively parallel processing) systems. A clear and flexible definition was needed to decide which of these systems was a supercomputer. This

definition needed to be independent of architecture. Because of Moore's Law, this definition also had to be dynamic to deal with the constant increase in computer performance.

Consequently, in early 1993 the TOP500 idea was developed by Prof. Hans Meuer and Erich Strohmaier at the University of Mannheim. The basic idea was to list the 500 most powerful computer systems installed around the globe and to call these systems supercomputers. The number 500 was picked based on our earlier market surveys, which indicated that more than 500 but fewer than 1,000 major vector systems had been installed at that time. The problem then was how to define how powerful a computer system is. For this task we decided to use the performance results of the Linpack benchmark from Jack Dongarra, as this was the only benchmark for which results were available for nearly all systems of interest [1].

Since 1993, we have published the TOP500 twice a year using Linpack results¹. Over the years the TOP500 has served well as a tool to track and analyze technological, architectural and other changes in the HPC arena [2]. Table 1 shows the top 10 systems as of June 2003. The TOP500 lists the Japanese Earth Simulator System as clearly the world's largest supercomputer since June 2002.

Table 1

Rank	Manufacturer	Computer	R _{max} [TF/s]	Installation Site	Country	Year	# Proc
1	NEC	Earth-Simulator	35.86	Earth Simulator Center	Japan	2002	5120
2	HP	ASCI Q, AlphaServer SC	13.88	Los Alamos National Laboratory	USA	2002	8192
3	Linux Networx/ Quadrics	MCR Cluster	7.63	Lawrence Livermore National Laboratory	USA	2002	2304
4	IBM	ASCI White SP Power3	7.3	Lawrence Livermore National Laboratory	USA	2000	8192
5	IBM	Seaborg SP Power 3	7.3	NERSC Lawrence Berkeley Nat. Lab.	USA	2002	6656
6	IBM/Quadrics	xSeries Cluster Xeon 2.4 GHz	6.59	Lawrence Livermore National Laboratory	USA	2003	1920
7	Fujitsu	PRIMEPOWER HPC2500	5.41	National Aerospace Laboratory of Japan	Japan	2002	2304
8	HP	rx2600 Itanium2 Cluster Quadrics	4.88	Pacific Northwest National Laboratory	USA	2003	1536
9	HP	AlphaServer SC ES45 1 GHz	4.46	Pittsburgh Supercomputing Center	USA	2001	3016
10	HP	AlphaServer SC ES45 1 GHz	3.98	Commissariat a l'Energie Atomique (CEA)	France	2001	2560

Table 1: Top 10 supercomputer systems as of June 2003.

¹ All data from the TOP500 and further analysis are available from our main web site at www.top500.org.

Performance Growth and Dynamic

One trend of major interest to the HPC community is the growth of the performance levels seen in the TOP500. Figure 1 shows the evolution of the total installed performance in the TOP500. We plot the performance of the first and last systems (at positions 1 and 500) on the list, as well as the total accumulated performance of all 500 systems. Fitting an exponential curve to the observed data points, we extrapolate out to the end of the decade. We see that our data validate the exponential growth of Moore's Law very well, even though we use Linpack performance numbers and not peak performance values. Based on the extrapolation from these fits we can expect to have the first 100 teraflop/s system by 2005. At that time, no system smaller than 1 Tflop/s should be able to make the TOP500 any more. Towards the end of the decade we can expect supercomputer systems to reach the performance level of 1 petaflop/s.

Figure 1




Figure 1: Performance growth in the TOP500 and extrapolation till the end of the decade.

The HPC market is by its very nature very dynamic. This is reflected not only by the coming and going of new manufacturers, but especially by the need to update and replace systems quite often to keep pace with the general performance increases. This dynamic is reflected in the average replacement rate of about 160 systems every half-year - or more than half the systems on the list every year. This means that a system which is at position 100 at a given time will fall off the TOP500 within two to three years.

Manufacturers

Now for a closer look at which companies actually produce the systems seen in the TOP500. In Figure 2 we see that 10 years ago, the specialized HPC companies such as Cray Research, Thinking Machines (TMC), Intel with their hypercube based iPSC systems, and the Japanese vector system manufacturers Fujitsu, NEC, and Hitachi dominated this market. This situation has clearly changed. Nowadays, mainstream computer manufacturers from the workstation and PC segment, such as IBM, Hewlett-

Packard, SGI and Sun, have largely taken their place. Cray, the last U.S. vector system manufacturer, is a notable exception and is now re-entering the market with the introduction of its new X1 computer system.

Figure 2




Figure 2: Manufacturers of systems in the TOP500.

System Architectures

The changing share of the different system architectures as reflected in the TOP500 is shown in Figure 3. Single-processor systems and SMPs with shared flat memory are no longer powerful enough to make the TOP500. For most of the last 10 years, MPP systems have dominated. During the last few years, the number of clustered systems grew considerably. Considering the impressive performance dominance of the vector-based Earth Simulator System, it is an interesting and open question as to what share of the TOP500 traditional supercomputers will be able to hold on to.

Changes in computer architecture also make it more and more of a challenge to achieve high performance efficiencies in the Linpack benchmark used to rank the 500 systems. With knowledge and effort, the Linpack benchmark can still be implemented in very efficient ways as recently demonstrated by a new implementation developed at the U.S. Department of Energy's National Energy Research Scientific Computing (NERSC) Center for their 6,656-processor IBM SP system.

Figure 3




Figure 3: Dominant supercomputer system architectures. Constellations (Const.) are cluster of large SMPs.

Processor Architectures

With respect to the processors used, the HPC market has always been different from the mainstream computing markets. The custom vector processors used in the '70s and '80s were replaced in the early '80s by a mix of custom RISC processors and later on – finally - by mainstream superscalar processors such as the IBM Power processor, MIPS processors, or HP PA-Risc processors. The most noticeable difference between HPC and the overall computer market is that for much of the last decade, systems based on the Intel microprocessor played only a minor role in the HPC arena, as shown in Figure 4. One reason for Intel's absence in this market is almost certainly due to the company's decision to abandon its HPC ambitions in the mid-1990s. The advent of PC clusters and their slow appearance in the TOP500 helped to increase the number of Intel-based supercomputers again, and as of June 2003, Intel is again a main provider of processors, along with HP and IBM, for TOP500-class systems.

Figure 4




Figure 4: Chip technology of systems in the TOP500.

Main Supercomputing Sites

Government programs such as the Department of Energy's ASCI (Advanced Simulation and Computing) program certainly attract a lot of public interest. It is not clear, however, to what extent these programs are actually capable of influencing the market directly in the short term as they only represent isolated (but large) business opportunities, which are still small compared to the overall market size. In the long term, U.S. government programs do certainly provide an environment for HPC system users and producers to establish, defend and increase their competitive advantage.

This can be seen by analyzing the combined 10-year history of the TOP500. The Linpack performance for a system in a specific TOP500 edition is "normalized" by showing the ratio of its Linpack performance to the sum of the Linpack performances for all the systems on that list. Defining normalized performance in this way removes the influence of Moore's Law and allows us to generate aggregate statistics over all 21 editions of the TOP500, giving equal weight to early lists.

For all the centers, we add up the hypothetical normalized Linpack performance that all of their systems could have delivered over their lifetime. The list of the top 10 centers assembled in this fashion is shown in Table 2. We see that there are seven centers from the United States, three from Japan and none from Europe. The first three centers are the ASCI centers. The other seven centers together provided roughly the same number of compute cycles as the three ASCI centers. The strong influence of government programs on very large centers is clearly evident.

Table 2

Site	sum % norm. Rmax	Sum Rmax TF/s	Country
1 Sandia National Laboratories	85.1%	37.21	US
2 Los Alamos National Laboratory	80.1%	63.82	US
3 Lawrence Livermore National Lab.	66.3%	94.24	US
4 NAL	47.7%	11.02	Japan
5 Earth Simulator Center	37.9%	107.58	Japan
6 University of Tokyo	34.6%	19.71	Japan
7 Pittsburgh Supercomputing Center	29.8%	23.84	US
8 NERSC/LBNL	29.0%	24.98	US
9 Oak Ridge National Laboratory	28.0%	15.95	US
10 NAVOCEANO	24.2%	20.53	US

Table 2: Top 10 centers determined by the sum of the normalized linpack performance for all systems installed in a center. Classified sites are excluded from this analysis.

The lack of comparable European programs is also reflected by the absence of any European center in this table. If this situation continues, European scientists might find themselves in a position of only having access to computer resources which are an order of magnitude smaller than in the USA.

Conclusion

The HPC market has always been dominated by very rapidly changing technologies and architectures. The speed of this change is ultimately coupled to Moore's Law, which states that computing capabilities grow by roughly a factor of 2 every 18 months. Tracing the evolution of such a dynamic marketplace is a challenge and the tools and methods used for this have to be re-evaluated constantly. This is no different for the TOP500 project. In 1993 we decided to switch from our old form of HPC market statistics to the TOP500 in its current form, and it has served us well since then. In the last 10 years the diversity of architectures and applications in the HPC market has increased substantially. It has to be kept in mind that doing justice to this large variety is certainly not possible with any single benchmark, and we are evaluating several approaches to improve this situation. This includes ongoing projects for the creation of new benchmarking metrics such as those developed by the Performance Evaluation Research Center in DOE's Scientific Discovery through Advanced Computing (SciDAC) program [3].

Bibliography

- 1 Dongarra, J., Performance of Various Computers Using Standard Linear Equations Software. 2000, University of Tennessee, Computer Science Tech Report CS-89-85: Knoxville.
- 2 Erich Strohmaier, Jack J. Dongarra, Hans-Werner Meuer, and Simon. Horst D. The Marketplace of HPC. Parallel Computing, 25th anniversary edition, North Holland, 25():1517-1544, 1999.
- 3 See <http://perc.nerisc.gov> for current details.